

# Brief Notes on Hard Takeoff, Value Alignment, and Coherent Extrapolated Volition

GOPAL P. SARMA\*

School of Medicine, Emory University

## Abstract

I make some basic observations about hard takeoff, value alignment, and coherent extrapolated volition, concepts which have been central in analyses of superintelligent AI systems.

## I. ON HARD TAKEOFF

The distinction between *hard takeoff* and *soft takeoff* has been used to describe different possible scenarios following the arrival of human-level artificial intelligence. The basic premise underlying these concepts is that software-based agents would have the ability to improve their own intelligence by analyzing and rewriting their source code, whereas biological organisms are significantly more restricted in their capacity for self-improvement [1, 2, 3, 4, 5].

There is no precise boundary between the two scenarios, but in broad strokes, a hard take refers to a transition from human level intelligence to superintelligence in a matter of minutes, hours or days. A soft takeoff refers to a scenario where this transition is much more gradual, perhaps taking many months or years. The practical importance of this qualitative distinction is that in a soft takeoff, there may opportunities for human intervention in the event that the initial AI systems have problematic design flaws.

The purpose of this brief note is simply to point out the following: takeoff speed refers to the rate of change of the agents' *level of intelligence*, and not our perceived changes in

the world around us. Because the notion of an "intelligence explosion" has been constructed in analogy to a physical explosion [6], it gives rise to an inaccurate mental picture in people's minds. If self-improving AI systems<sup>1</sup> are thought to be the intellectual analogue of nuclear chain reactions, then the natural image of an intelligence explosion that this metaphor creates is a scenario in which massive, disruptive changes take place in the world that are difficult for individuals and for society to handle.

However, the premise of intelligent agents with capacities in substantial excess of any human being, which are able to process the sum total of human knowledge in the form of books, video, and ongoing contemporary events implies that greater levels of intelligence in the AI systems will be accompanied by actions taken with a corresponding level of information, insight, and operational skill. Therefore, if the initial systems are designed correctly with respect to value alignment and goal structure stability, it is in fact a hard takeoff scenario which would be *less disruptive*

---

<sup>1</sup>As this analysis neither requires nor implies that the driving force of change is a unitary agent, I have chosen to use the plural terms "software-based agents," "AI systems," or "superintelligent AI systems." It may very well be a collection of agents / systems possessing powerful AI capabilities in aggregate.

---

\*Email: gopal.sarma@emory.edu

than a soft takeoff, not the other way around.

I reiterate this claim for emphasis: *The take-off speed of an intelligence explosion refers to the rate of change of intelligence in the AI systems, and not our perceived changes in the world around us. Therefore, under the assumption of correctly design systems, a hard takeoff is preferable to a soft takeoff because the resultant changes that take place in the world will be executed with greater precision, thoughtfulness, and insight.*

## II. ON VALUE ALIGNMENT AND COHERENT EXTRAPOLATED VOLITION

The preceding argument relied on a key assumption, namely that the AI systems capable of self-improvement were designed correctly with respect to value alignment and goal structure stability. Value alignment refers to the construction of systems that take actions consistent with human values. Russell states 3 design principles which encapsulate the notion of value alignment [7]:

1. The machine's purpose must be to maximize the realization of human values. In particular, it has no purpose of its own and no innate desire to protect itself.
2. The machine must be initially uncertain about what those human values are. The machine may learn more about human values as it goes along, but it may never achieve complete certainty.
3. The machine must be able to learn about human values by observing the choices that we humans make.

A related notion is Yudkowsky's "coherent extrapolated volition" [8, 9]. The basic premise of this proposal is that sophisticated AI systems will be capable of extrapolating and resolving disagreements between the value systems of individuals and groups, ultimately arriving at a goal structure that represents the collective desires of humanity. This process of iterated reflection is analogous to Rawls'

"reflective equilibrium" [10].

Like the metaphor of hard takeoff, the notions of value alignment and coherent extrapolated volition can also give rise to an inaccurate mental picture, namely, that the aligned goal structure would either require or result in all humans arriving at complete agreement on all issues. However, with adequate resources, it may very well be that value-aligned AI systems shape a world in which groups of individuals co-exist who disagree about object-level issues. Certainly we can point to many examples in contemporary human society where individuals maintain divergent preferences without conflict.

The purpose of this brief note is simply to point out the following: *Implicit in any practical analysis of value alignment are the physical resources available to the AI systems. In particular, the construction of a human compatible goal structure does not mean that all human disagreements have been resolved. Rather, it means that a mutually satisfactory set of outcomes has been achieved, subject to resource constraints.*

It may be impossible to arrive at a consensus goal structure without adequate resources. As a trivial example, if we have two individuals each of whom desires at least one apple, there is no disagreement if we have two apples. On the other hand, if there is only one apple, conflict may very well be inevitable in the absence of other factors with which to resolve the imbalance between specific individual desires and total available resources. In the context of superintelligent AI systems capable of exerting substantial influence on the world and shaping society on a global scale, adequate analysis of value alignment requires an understanding of the sum total of physical resources that the AI systems have at their disposal.

## Acknowledgements

I would like to thank Eric Drexler and Anders Sandberg for valuable discussions and feed-

back on the manuscript.

## REFERENCES

- [1] N. Bostrom, *Superintelligence: Paths, Dangers, Strategies*. OUP Oxford, 2014.
- [2] I. J. Good, "Speculations Concerning the First Ultraintelligent Machine," *Advances In Computers*, vol. 6, no. 99, pp. 31–83, 1965.
- [3] D. Chalmers, "The Singularity: A Philosophical Analysis," *Journal of Consciousness Studies*, vol. 17, no. 9-10, pp. 7–65, 2010.
- [4] M. Shanahan, *The Technological Singularity*. MIT Press, 2015.
- [5] C. Shulman and A. Sandberg, "Implications of a software-limited singularity," in *Proceedings of the European Conference of Computing and Philosophy*, 2010.
- [6] E. Yudkowsky, "Artificial intelligence as a positive and negative factor in global risks," in *Global Catastrophic Risks* (N. Bostrom and M. Cirkovic, eds.), ch. 15, pp. 308–345, Oxford: Oxford University Press, 2011.
- [7] S. Russell, "Should We Fear Supersmart Robots?," *Scientific American*, vol. 314, no. 6, pp. 58–59, 2016.
- [8] E. Yudkowsky, "Coherent Extrapolated Volition," *Singularity Institute for Artificial Intelligence*, 2004.
- [9] N. Tarleton, "Coherent extrapolated volition: A meta-level approach to machine ethics," *Machine Intelligence Research Institute*, 2010.
- [10] J. Rawls, *A Theory of Justice*. Belknap, 1971.